Volume 10 Issue 02 Published, May 15, 2006 ISSN 1535-864X DOI: 10.1535/itj.1002

Intel® Technology Journal

Intel® Centrino® Duo Mobile Technology

System Memory Power and Thermal Management in Platforms Built on Intel® Centrino® Duo Mobile Technology

More information, including current and past issues of Intel Technology Journal, can be found at:
http://developer.intel.com/technology/itj/index.htm

System Memory Power and Thermal Management in Platforms Built on Intel® Centrino® Duo Mobile Technology

Jayesh Iyer, Mobility Group, Intel Corporation
Corinne L. Hall, Mobility Group, Intel Corporation
Jerry Shi, Mobility Group, Intel Corporation
Yuchen Huang, Technology and Manufacturing Group, Intel Corporation

Index words: TS on DIMM, DT in SPD, throttle, thermal, memory performance

ABSTRACT

Specific form factor requirements of the mobile thin and light platform limit the kinds of cooling solutions these platforms can have. As a result, individual components inside mobile platforms (such as system memory devices) can heat up. With increased memory speeds and capacities, we are now reaching the point where the memory thermals are starting to exceed the cooling capabilities of mobile systems. To ensure that the memory devices operate within their thermal limits, their case temperatures need to be monitored and memory accesses throttled if any memory device overheats. In this paper, we discuss the need for memory throttling and address two memory throttling techniques, implemented in platforms built on Intel® Centrino® Duo mobile technology. These techniques greatly improve the memory power/thermal management in a thermally constrained platform by maximizing performance while keeping the system within its thermal limits.

First, we describe Delta Temperature (DT) in Serial Presence Detect (SPD), a novel memory throttling technique, which uses the Virtual Temperature Sensor (VTS) throttling mechanism (VTS was first implemented in the Intel® 815GMCH chipset) to optimize memory throttling control. VTS provides the means to predict the temperature of the memory devices in platforms that don't have a physical thermal sensor on the DRAM modules. When the predicted temperature exceeds the set memory thermal limit, throttling is enforced. We also discuss the implementation details of DT in SPD and the performance benefits it offers in terms of guardband reduction and bandwidth recovery, over the existing open loop throttling techniques.

Finally, we focus on Thermal Sensor (TS) on a Dual In-line Memory Module (DIMM), a closed loop solution for memory throttling control. This technique uses a physical thermal sensor on each of the memory modules that signals the chipset to throttle memory traffic when the memory module exceeds the thermal trip-point. Adding a thermal sensor to the memory module allows the system to continue running at full bandwidth until the critical temperature is reached. In addition to providing both a mechanism for preventing the DRAMs from exceeding the maximum temperature specification (85°C for memory) and a mechanism to control skin temperature, we discuss the performance benefits that TS on DIMM offers over existing VTS-based solutions, including guardband reduction and bandwidth recovery.

INTRODUCTION

Driven by demand and ever more efficient technological advancements, as the form factors of computing platforms get smaller and component densities get higher, system power consumption and thermal issues have become more challenging than ever before, especially in the case of laptops, slim desktops, and blade servers. Naturally, each system has a total cooling capacity per its specific design, and each major component has a cooling limit or power consumption allowance for balancing the system performance in that design. When the total system cooling capacity is smaller than the sum of each component's Thermal Design Power (TDP), as seen in many small-form factor systems, all the components simply cannot simultaneously operate at their TDP level since the system cannot cool them. In such a system, it is also unacceptable to allow one component to free-run without power consumption and thermal restrictions as that would leave the performance of other components to suffer. Balanced cooling limits for major components should thus be achieved at the design level. Having briefly laid the background on cooling limits, we now focus on system memory (specifically in laptops), which is a major power-consuming component of platforms; specifically we

discuss its cooling limits and enhancement in performance through better power/thermal management (memory bandwidth recovery using new throttling techniques).

To better understand the variation in memory cooling limits from system to system, a glimpse at three popular categories of laptops is helpful. The first of these categories is the Thin & Light (T&L) laptop, which is the mainstream, more conventional model. These laptops have a Z-height of around 1.1''-1.2'' and a memory cooling limit of 4-5 watts. The second category is the Mini-Note PCs, which have a smaller form factor than the T&L and a Z-height of around 0.9". They have a smaller cooling fan and a memory cooling limit of 2-2.5 watts. Finally there is the Sub-Note PCs, which have a form factor even smaller than the Mini-Note PCs. These laptops do not have a cooling fan and their cooling limit is around 1 watt. These data clearly show that mobile platforms have limited cooling capabilities, and as the form factors continue to get smaller, the cooling budget also shrinks considerably. With increased memory speeds and capacities, we are now reaching the point where the memory thermals are starting to exceed the cooling capabilities of mobile systems. When the cooling budget is exceeded, that means the system is no longer able to cool the memory subsection, and the DRAM case temperatures begin to exceed their maximum case temperature specification of 85°C. Our lab data, taken on multiple notebook systems while analyzing multiple Small Outline-Dual In-line Memory Modules (SO-DIMMs)[1], show that some 1 GB and greater capacity SO-DIMMs are exceeding their maximum specified case temperature when running a realistic workload at an ambient temperature of 35°C. Thus the memory bus needs to be throttled[2] to ensure that the DRAM devices operate within their thermal limits, reducing the risk of memory corruption and system instability.

Platforms built on Intel Centrino Duo mobile technology implement two memory throttling techniques to address this issue.

DELTA TEMPERATURE (DT) IN SERIAL PRESENCE DETECT (SPD)

Delta Temperature (DT) in Serial Presence Detect (SPD) [3] is a throttling technique that is particularly beneficial in platforms that do not have a physical thermal sensor on the memory modules. At this point in time, not all memory modules have a physical thermal sensor on the DRAM modules monitoring their case temperatures; therefore, their temperature needed to be tracked by some alternative means. A Virtual Temperature Sensor (VTS) throttling mechanism [1] implemented in the chipset (first implemented in the Mobile Intel® 915 Chipset family), provides this alternative solution. VTS predicts memory case temperature based on the measured memory access type at each clock cycle and the speed and type of memory in the system. DT in SPD provides DRAM thermal data for each memory access type (Reads, Writes, Self Refresh, etc.) and includes the DRAM maximum case temperature limit data in the SPD on the memory module. These data are used by the VTS throttle mechanism (which analyzes memory traffic) to estimate the DRAM case temperature and to throttle memory based on the thermal prediction. DT in SPD better predicts the power and thermal of the memory module and hence achieves greater guardband reduction over the open-loop throttling methods followed in the previous-generation mobile platforms, and thus enhances performance in terms of bandwidth recovery. Our experiments clearly demonstrate a performance benefit of almost 15-30% (in terms of sustained bandwidth) by using DT in SPD.

Though DT in SPD has its benefits in platforms that don't have a physical thermal sensor on the memory module, it is still an open-loop throttling technique since it does not actually know the operating temperature of the DIMM. Therefore it must always assume the laptop is operating under the maximum-allowed-room ambient temperature, which is typically considered to be 35°C for mobile environments. So if a notebook is actually running in an environment where the room ambient is only 20°C, then DT in SPD actually assumes it is running at 35°C and will initiate throttling ~15°C sooner than it needs to, thereby losing potential performance that could have been obtained had the system not initiated throttling so soon.

Thermal Sensor (TS) on Dual In-line Memory Module (DIMM)

TS on DIMM is the throttling technique that prevents memory from being over throttled and improves system performance. TS on DIMM integrates a physical thermal sensor onto the memory modules, providing real-time temperature information back to the chipset. The temperature feedback provides a closed-loop throttling mechanism that allows the system to adapt memory throttling to the actual environment the laptop is running in. So instead of just looking at memory traffic and estimating what the temperature of the DIMM *might* be based on estimated memory module power/thermals, now the chipset can wait to throttle until an actual thermal issue

[1] S0-DIMM: Small Outline Dual In-Line Memory Module (DRAM memory modules used in mobile platforms).

[2] Throttling: Solution to reduce memory traffic allowed on the memory bus.

[3] SPD stands for Serial Presence Detect. It's a Joint Electronic Device Engineering Council (JEDEC) spec.

arises. This allows systems running low-power DIMMs, or systems with excellent cooling mechanisms, or systems running at cool room ambient temperatures, to run high-bandwidth applications and not have to throttle memory until a critical temperature is reached. So unlike DT in SPD, TS on DIMM does *not* have to assume the worst-case room ambient temperature. This allows TS on DIMM to greatly reduce the guardband required and thus improve system performance by allowing the system to run unconstrained for longer periods of time. Our experiments clearly demonstrate, due to the guardband reduction achieved using TS on DIMM, the total bandwidth recovered is as much as 30% of the theoretical maximum bandwidth (running a particular application on a specific system and memory configuration).

In the following sections, we discuss the need for system memory throttling and describe the concept, system implementation, and benefits of the two above-mentioned throttling techniques: DT in SPD and TS on DIMM.

NEED FOR SYSTEM MEMORY THROTTLING

In this section we focus on the necessity for throttling the memory bus. The material forms the basis for our discussions in the rest of the paper.

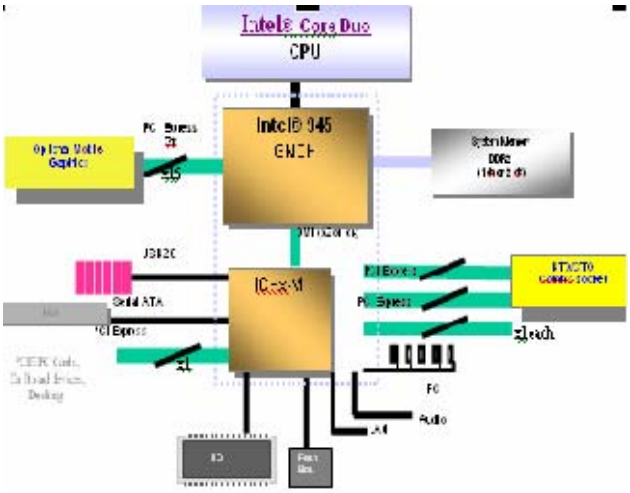

Figure 1: Schematic showing main components in current-generation laptops

Figure 1 shows a high-level view of the main components on a mobile platform built on Intel Centrino Duo mobile technology. The system memory gets accessed in almost all activities taking place in the platform. All data transfers to and from the system memory are managed by the Intel® chipset. Thus, if the chipset/system memory is idle, the platform itself is generally in an idle state. But if there is activity on the chipset/system memory, the platform is consuming more power, causing the chipset and memory components to heat up.

As mentioned earlier, mobile platforms have limited cooling capabilities. In the past, memory speeds and capacities generally allowed the system memory subsection to stay within the cooling limits of the platform, and the DRAM case temperatures were below the maximum operating specifications of 85°C. But with the increase in memory capacity and speed, we are now reaching the point where memory thermals are starting to exceed the cooling capabilities of mobile systems. When the cooling budget is exceeded, that means the system is no longer able to cool the memory subsection, and DRAM case temperatures begin to exceed their maximum case temperature specification.

In a recent laboratory study, several different thin-and-light laptop designs were tested with various SO-DIMMS of different memory speeds, capacities, densities, and vendors, using OpenGL Benchmark software. This software has high bandwidth utilization (about 52 percent of theoretical max) and causes DRAM devices to draw constant high power.

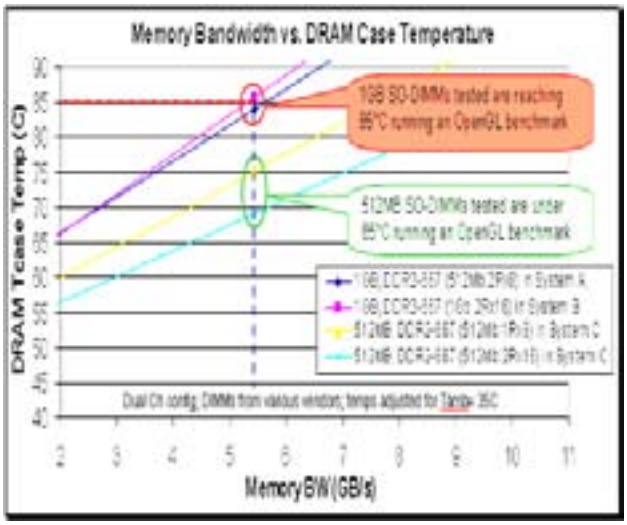

Figure 2: Lab data showing memory bandwidth and DRAM case temperatures

In Figure 2, temperature and bandwidth data are dependent on system, memory, and software configuration. The figure shows the results of a small sample of this thermal study. All three systems are thin and light designs: each system is from a different vendor. System A platform layout has one memory module on top and one on the bottom, System B platform layout has both memory modules on top, and System C platform has both memory modules on the bottom.

The results of the above study show that typical 512 MB SO-DIMMs in various thin-and-light notebook designs are

well below 85°C and do not appear to be in jeopardy of exceeding their thermal limits. In some cases, however, 1 GB-capacity SO-DIMMs are reaching, and sometimes exceeding, their maximum case temperature specification of 85°C. These results show that memory modules are already operating near their maximum specifications, and as memory power and thermals increase with system capacity and speed, memory modules will begin exceeding their maximum specification with multiple realistic workloads. Small form factor designs are of even greater concern due to their lower cooling budgets and thermally challenged environments. There is clearly a need for a robust thermal management solution.

If left unchecked, DRAM devices will start running above their maximum operating case temperatures, and memory-related reliability issues will begin to crop up. Thus the memory bus needs to be throttled to ensure that the DRAM devices operate within their thermal limits. Memory throttling provides a solution to cool the DRAM devices by reducing memory traffic allowed on the memory bus, thereby reducing the power consumed by the DRAM devices and thus reducing thermal output.

We now describe the two throttling techniques in detail.

DT in SPD

DT in SPD is a throttling technique that gives a huge performance benefit to platforms that do not have a physical thermal sensor on the memory module. Before going into the details of DT in SPD, let's briefly look into the throttling methodology followed in the chipsets of previous-generation platforms.

A DDRx[4]-based memory subsystem, as a major power-consuming component on a platform, is peculiar in that its power consumption is difficult to predict in reality. The reason for this is that memory vendors have their own unique processes and design technologies. This results in large variations among different DRAM vendors in power consumption given a fixed device density and speed for each particular DDRx technology. Figure 3 shows the differences in the Burst read current (IDD4R) of the best-, typical-, and worst-case memory supplier (DDR2 512 MB DRAM) across three different memory speeds. As the memory device power is a function of its IDD numbers, a huge difference in the IDD numbers between the best-case and worst-case memory suppliers translates into a huge difference in their power consumption.

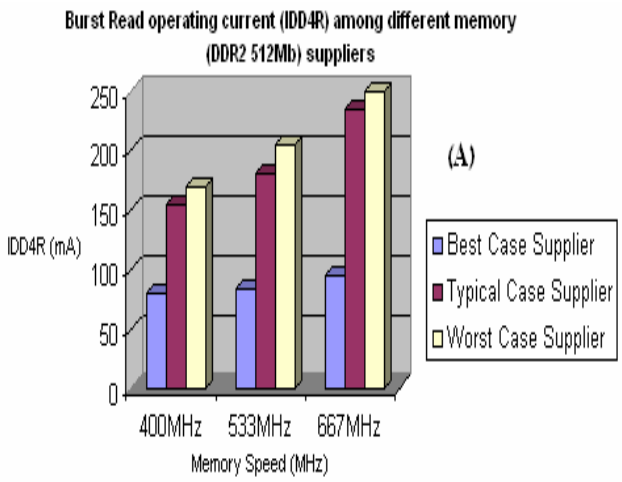

Figure 3: Power differences between best-, typical- and worst-case memory suppliers

In the previous-generation mobile platforms, the system could tell the memory type (device density, number of ranks, device width) and its vendor by retrieving the data from a small EEPROM (called SPD) on each memory module during boot up. But this doesn't indicate how much power the module consumes. Also these platforms did not have a thermal sensor on the DRAM modules to monitor their temperature, and hence their temperature needed to be tracked by some alternative means.

The VTS throttle mechanism (implemented for the first time in the Mobile Intel 915GMCH Express Chipset) provides this alternative solution. VTS predicts the memory case temperature based on the memory access type at each clock cycle. The power estimates for each of the memory access types are programmed in the chipsets throttle weights register [2] and a discretized lumped thermal capacitance model [6] is used to predict the time-varying temperature response to the power consumption. The chipsets in these platforms implemented a throttling methodology without any feedback on power or temperature from the memory module. Without this knowledge of power consumption from the memory module, the system (VTS throttle mechanism implemented in the chipset) was forced to assume the worst-case power consumption for design safety. Thus memory thermal limits (throttling threshold value) were set based on the worst-case supplier's power data. This approach thus resulted in over-guard-banding[5] and over-memory-throttling. The memory performance degrades as low-power memory modules are treated as the worst-case modules.

[4] X is 1, 2 or 3 based on the Double Data Rate Synchronous DRAM memory technology (DDR, DDR2, DDR3)

[5] Reduction in bandwidth to ensure power/thermal values stay within design safety limit.

Now the inevitable question is, how can a system, especially when it does not have a physical thermal sensor on the memory modules, accurately and cost-effectively tell memory power consumption to avoid over throttling memory, i.e., how can it reduce the guardband? It can by using DT in SPD.

Overview of DT in SPD

DT in SPD is a power/thermal prediction scheme providing DRAM power/thermal data in the SPD on the module. It stores key temperature rise data for each memory access type (Reads, Writes, Self-Refresh, etc.) and DRAM maximum case temperature limit data (Tcasemax) in the SPD on the memory module. The system (VTS throttling mechanism implemented in the chipset) uses this information to better estimate the temperature of the DRAM devices and to determine when throttling is necessary. When process shrinks or other power optimizations occur and DRAM power dissipation decreases, the system uses the information stored in SPD to reduce memory throttling and regain system performance.

Usage Model

Memory module vendors report the delta temperature rise parameters and Tcasemax in SPD. These parameters are read by the BIOS from the SPD at boot time. Subsequently, the system adjusts memory throttle limits based on these parameters.

Performance Benefit of DT in SPD

DT in SPD greatly reduces the guardband and helps recover the bandwidth as compared to the throttling methodologies implemented in chipsets in the previous platforms, which set memory thermal limits (throttling threshold value) based on the worst-case supplier's power data. A lab study was done on thin and light laptops using memory modules of different configurations from different vendors. As mentioned earlier, there is a huge difference in the power numbers between different vendors (for the same configuration). The best-case, typical-case, and worst-case DRAM vendors (for each configuration of the memory) in terms of power numbers were used for the study. A chipset stress utility based on Windows*, which generates very high traffic (almost all Page hits) on the memory bus, was used to exercise the system memory. Figures 4 and 5 each give the comparison of performance achieved running the utility with 30% Writes and 70% Reads on the thin and light design. It shows the sustained bandwidth across different memory throttle limits using the two throttling control methodologies: 1) with DT in SPD; 2) without DT in SPD (using worst-case supplier power data).

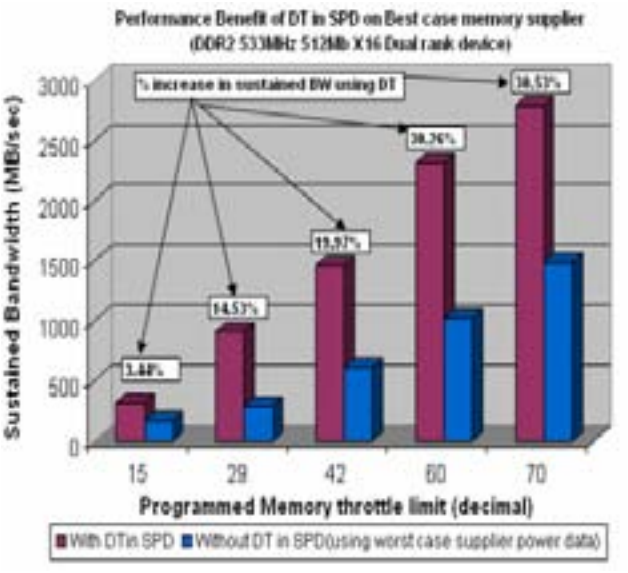

Figure 4: Sustained bandwidth across memory throttle limits (best-case supplier)

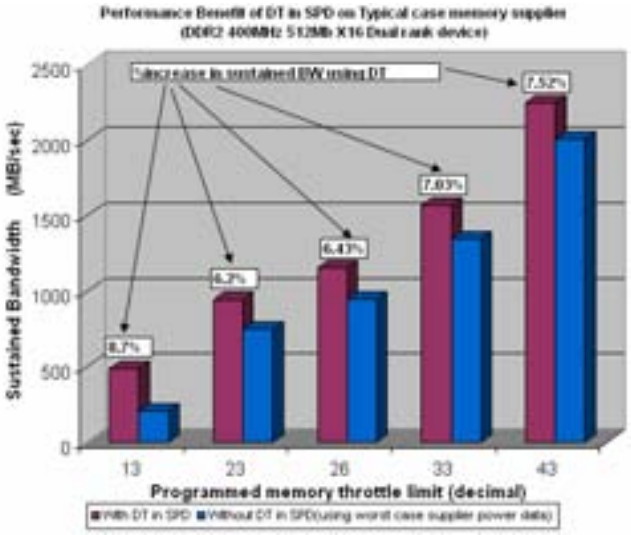

Figure 5: Sustained bandwidth across memory throttle limits (typical-case supplier)

The results in Figure 4 show that a typical-case memory supplier gives a performance benefit of 8% (in terms of sustained bandwidth) using DT in SPD. But for measurements done on a best-case memory supplier, the performance benefit goes up to 30% (in terms of sustained bandwidth) using DT in SPD as shown in Figure 5. Thus the results clearly demonstrate that using DT in SPD greatly reduces the guardband and helps recover the bandwidth. This enhances memory performance, especially for very low-power memory modules (best-case supplier).

Though DT in SPD has its benefits in systems that don't have a physical thermal sensor on the memory modules, it

is still an open-loop throttling mechanism as it does not actually know the operating temperature of the DIMM. Therefore it must always assume the notebook is operating under the maximum allowed room ambient temperature, which is typically considered to be 35°C for mobile environments. So if a notebook is actually running in an environment where the room ambient temperature is only 20°C, then DT in SPD actually assumes it is running at 35°C and will initiate throttling ~15°C sooner than it is needed, thereby losing potential performance that could have been obtained had the system not initiated throttling so soon.

TS on DIMM provides a much more robust thermal throttling mechanism while optimizing performance by offering reduced guardband over other existing methods. The following section describes in detail the concept, implementation details, and the performance benefit achieved using TS on DIMM.

TS ON DIMM

TS on DIMM is a closed-loop throttling technique that offers a more advanced approach to system thermal management through the use of a physical thermal sensor to further reduce guardband present in the existing methods. The reduction of guardband has a direct and positive impact on system performance.

Overview of TS on DIMM

Instead of using power *prediction* to estimate the temperature of the DRAM case, TS on DIMM uses a physical thermal sensor integrated on the DIMM module to monitor the temperature of DIMM. If the thermal sensor detects that the DIMM temperature is exceeding a programmable critical trip point, it triggers an event signal that tells the Graphics Memory Controller Hub (GMCH) to throttle the memory traffic, thereby reducing the DRAM case temperature. The addition of a physical thermal sensor results in a closed-loop throttling methodology that allows for real-time throttling based on the measured temperature. This closed-loop methodology leads to reduced guardband primarily due to the ability to sense ambient temperature. Most of the guardband present in existing throttling mechanisms is due to their open-loop policy, which means that since they do not receive feedback on the actual ambient or DIMM temperatures, they must always assume the worst-case scenario. TS on DIMM provides the needed temperature feedback and allows the system to hold off on throttling until the actual DIMM temperature reaches a programmed critical temperature trip point. The closed-loop policy offered by TS on DIMM is what allows it to reduce the guardband over existing memory throttling methods.

System Implementation of TS on DIMM

System implementation for TS on DIMM is very straight forward; it requires minor changes to the platform, the memory module, and BIOS. Figure 6 provides a block diagram showing the components involved and the logical connections required for TS on DIMM support. The thermal sensor [3] located around the central area of a memory module is connected to the chipset via three wires, two of them for SMBus[6], and the third one for the Event# signal. TS on DIMM implementation does not require any involvement with signals on the memory bus. Any time the temperature of the sensor exceeds the threshold limit programmed by BIOS, the thermal sensor asserts the Event# signal that triggers the chipset to initiate memory throttling until Event# de-asserts. The thermal sensor will de-assert the Event# pin after the temperature goes below the critical trip point.

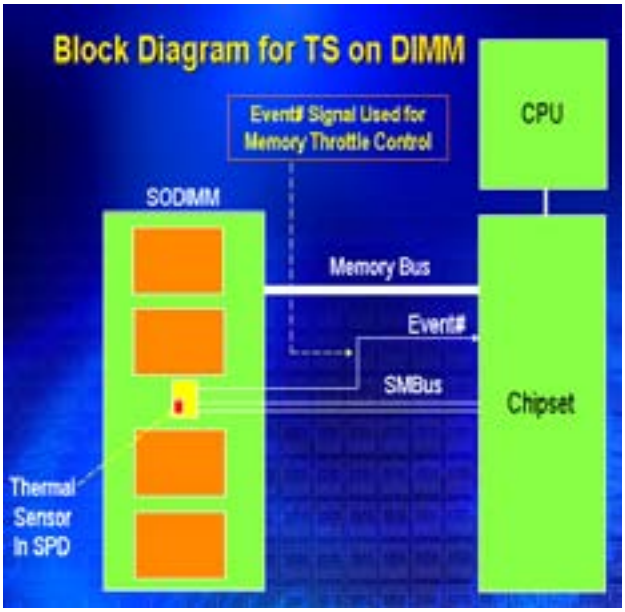

Figure 6: Block diagram for TS on DIMM system implementation

The SMBus master is inside the chipset. The sensor acts as a slave sitting on the SMBus with SPD on a module. Physically, the sensor and SPD [4] can be either integrated into one package as shown in Figure 3 or separated as two stand-alone parts on a module. However, the thermal sensor and the SPD are two separate logical functions using two different SMBus addresses. At power up, BIOS read data from the SPD and initiate the thermal sensor by programming its registers via SMBus, including setting the critical threshold value. When the initiation is

[6] System Management Bus defined by Intel for low-speed system management communication.

complete, the thermal sensor will begin monitoring the temperatures and assert the Event# signal when the temperature of the sensor exceeds the critical trip point. Once the thermal sensor is initiated the system can also poll the thermal sensor via the SMBus at any time and monitor the temperatures real-time.

Platform Requirements to Support TS on DIMM

Routing and hardware support at the platform level is simple; they just require routing the Event# signal from the GMCH [5] to each SO-DIMM connector and placing a 10k ohm pull-up resistor on this signal. For Intel chipsets, the Event# pin functionality and routing is supported on mobile platforms starting from 2006. DDR2 and DDR3 industry-standard SO-DIMM connectors both support TS on DIMM, which only required one additional connection for the Event# signal. SMBus signals were already routed through the connector for the SPD. From a platform perspective, that is all that is required to support TS on DIMM.

Thermal Sensors and Memory Modules

Thermal sensors for memory modules come in two varieties: remote and integrated. Remote thermal sensors are stand-alone 8-pin thermal sensors that can be placed on the memory module and used for monitoring temperatures. This remote thermal sensor option would be used in conjunction with standard SPD EEPROM devices used on modules today. The second option is to integrate the thermal sensor feature into the existing SPD device, also called TS in SPD, thus removing the need for an additional discrete part. Both versions operate in the same manner; they share the SMBus connections as well as the Event# signal connection. So from a platform routing perspective there is no difference between the two implementations: the module could have support for both, although only one implementation is used at a time.

Remote thermal sensors are available in the market today for use in memory module applications. The integrated TS in SPD devices are still in development with some samples starting to be available Q2'06.

Select DDR2 memory modules from limited suppliers will support TS on DIMM; some will support the remote thermal sensor implementation while others will support the integrated option.

RESULTS

Performance Benefits of TS on DIMM

The reduction in guardband that TS on DIMM offers over existing methods gives it a performance edge in both bandwidth as well as benchmark scores when running high-bandwidth applications. For example, when using an OpenGL benchmark on 1 GB DDR2-667 SO-DIMMs in dual channel mode, lab data showed that for every degree C of guardband removed, the system saw up to 220 MB/s bandwidth improvement per degree C, as shown in Figure 7. Since TS on DIMM typically saves ~15°C of guardband over existing methods when running a high-bandwidth application at room temperature, and at 220 MB/s per degree C, the total bandwidth recovered with TS on DIMM can be approximately 3 GB/s for this particular application, which is **30% of the theoretical maximum bandwidth** for this system and memory configuration. This bandwidth recovery assumes that the workload running on the system could utilize more bandwidth if allowed, and the bandwidth recovered is also dependent on the application used and the system configuration.

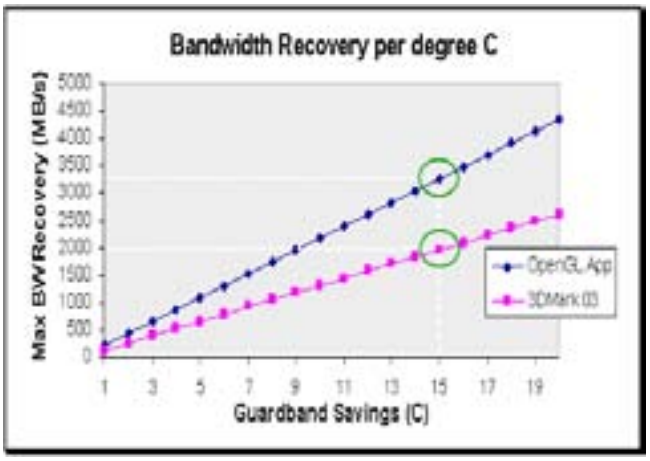

Figure 7: Bandwidth recovered per degree of guardband removed

Due to the increased bandwidth per degree C of guardband savings, benchmark scores also show improvements with TS on DIMM, as shown in Figure 8. (Benchmark score recovery is dependent on system, memory, and software configuration.) Using the same 1 GB, DDR2-667 SO-DIMMs in dual channel mode as described above, for every degree C of guardband removed, the system saw ~3.3 frames per second per degree C for a particular OpenGL application, and up to 13 3DMarks score recovery with 3DMark 03 per degree C. Again, since TS on DIMM typically saves ~15°C over existing methods, that results in an approximately 50 frames per second improvement as well as about a 200 3DMarks score improvement with these applications. Results may vary with different applications and different system configurations.

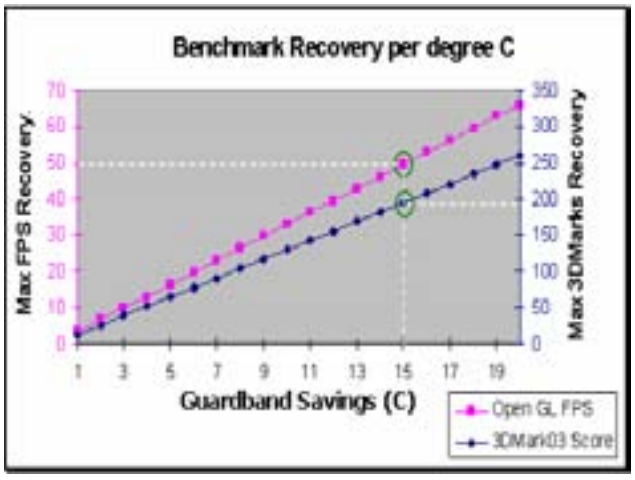

Figure 8: Benchmark scores recovered per degree of guardband removed

To further understand how the reduction in guardband impacts bandwidth availability, let's look at the thermal study from earlier in this paper and determine where TS on DIMM would throttle the worst-case 1 GB SO-DIMM module in comparison to DT in SPD. Please see Figure 9 below (temperature and bandwidth data is dependent on system, memory, and software configuration).

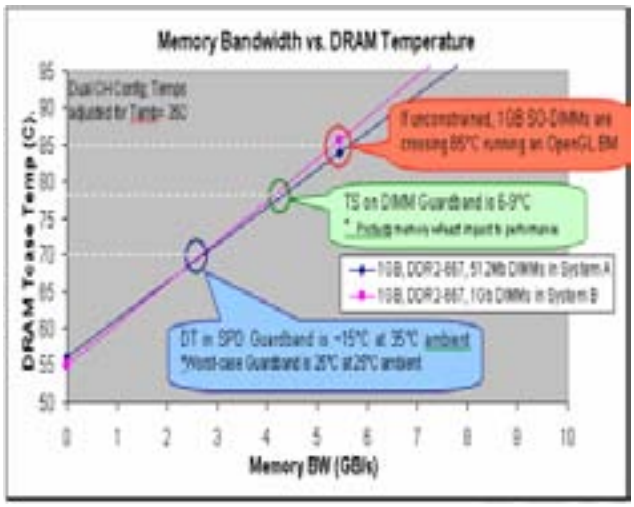

Figure 9: Guardband comparison between TS on DIMM and DT in SPD

The graph above indicates DT in SPD would throttle when the DRAM temperature reaches 70°C (bandwidth available is 2.5 GB/s) while TS on DIMM would throttle only at 79°C (enabling memory bandwidth of 4.5 GB/s). The results clearly show that the reduction in guardband directly impacts bandwidth availability as the temperature of the DIMM reaches and exceeds 60°C at a room ambient temperature of 25°C (or 70°C if the ambient temperature is 35°C).

Other Benefits of TS on DIMM

Aside from the performance benefits mentioned above, TS on DIMM offers other benefits over existing methods. First, since the thermal sensor is measuring true operating temperature, it is able to help protect the DRAM cases from exceeding their maximum specifications even if an outside source is heating the DIMMs. For example, if a fan breaks or if the notebook is sitting in direct sunlight, the thermal sensor is always monitoring the DRAM temperatures and if they exceed the critical trip point, the chipset will throttle the memory traffic and protect the DRAMs.

With the existing methods of temperature prediction based on memory traffic alone, the chipset has no idea if some abnormality is causing the DRAM temperatures to rise beyond their specifications. Since TS on DIMM is monitoring the temperature of the memory, it also provides a way for controlling notebook skin temperatures. Keeping memory within certain thermal limits will also help keep skin temperatures within certain thermal limits.

SUMMARY

Specific form factor requirements of notebooks limit them from having exhaustive cooling solutions. As a result, individual components inside the notebook (such as system memory devices) can heat up to the point of exceeding their operating specifications. To ensure that the memory devices operate within their thermal limits, their case temperature needs to be monitored and memory accesses throttled if any memory devices approach their thermal limits. Current lab data taken on multiple notebooks and from analyzing multiple SO-DIMMs show that 1 GB capacity SO-DIMMs and greater are nearing their maximum specified case temperature when running a realistic workload at an ambient temperature of 35°C. The data were taken from Thin and Light notebook designs: the thermal concern for small form factor designs are of even greater concern due do their limited cooling abilities.

Open-loop throttling mechanisms, including DT in SPD, throttle memory based on a thermal prediction scheme that analyzes memory traffic to estimate the temperature of the DRAM case temperatures. These mechanisms lead to over-guardbanding since they always have to assume worst-case environment conditions, such as the ambient temperature of the room. Also, because these mechanisms are estimating thermal parameters only based on memory traffic, they cannot protect all DRAMs from exceeding their maximum case specifications of 85°C.

TS on DIMM, a closed-loop throttling mechanism, offers several benefits over open-loop throttling methods. The

closed-loop methodology of TS on DIMM allows it to reduce the guardband substantially over existing methods, increasing system performance by allowing the system to run unconstrained longer before initiating throttling. TS on DIMM also provides a way to control notebook skin temperatures and provides a safeguard mechanism to help prevent the DRAMs from exceeding their maximum case specification of 85°C.

Currently TS on DIMM is an optional feature for DDR2 SO-DIMMs. Intel is working with memory industry to drive this feature into DDR3.

ACKNOWLEDGMENTS

Thanks to David Wyatt for writing the initial remote thermal sensor spec and doing the initial study that enabled all future work. We thank Rajeev Menath for his support during TS on DIMM measurements, and Ishmael Santos and Rafael Rodarte for helping during the test plan development as well as supporting the sensitivity study. Thanks to Christopher E. Cox and Rachael Young for helping drive this feature in JEDEC. Thanks to Bill Sanderson for reviewing the results and Deepa Mohan for reviewing the paper and providing feedback. Also, thanks to Ramesh Shetty for his immense help during the measurement setup in the lab. Special thanks to Nilesh Shah and Minh Ngyuen for driving this feature.

REFERENCES

[1] Eric Samson, Sridhar V. Machiroutu, Je-Young Chang, Ishmael Santos, Jim Hermerding, Ashay Dani, Ravi Prasher, Dawid Song, "Interface Material Selection and Thermal Management Technique in Second-Generation Platforms Built on Intel Centrino Mobile Technology," in *Intel Technology Journal*, Volume 09, Issue 01, February 2005.

[2] Jayesh Iyer, Yuchen Huang, Jerry Shi, "DT in SPD: A Better Approach to System Memory Throttling" at *Intel Design & Test Technology Conference*, August 15-19, 2005, Portland, OR.

[3] "Mobile Platform Memory Module Thermal Sensor Component Specification," *JC-42.4*, Ballot proposal (Item number: 1640.07).

[4] JESD21C JEDEC DDR2 SPD Revision 1.1.

[5] "RS-Mobile Intel® 955XM 945 GM/PM/GMS and 940 GML Express Chipset," *External Design Specification* (EDS), Vol. 1 & 2.

[6] R. S. Prasher, J. Shipley, S. Prstic, P. Koning, and J. L. Wang, "Thermal Resistance of Particle Laden Polymeric Thermal Interface Materials," in *Proceedings of International Mechanical Engineering Congress and Exposition*, Nov. 15-21, 2003, Washington, D.C.

AUTHORS' BIOGRAPHIES

Jayesh Iyer joined Intel's Mobile Platforms group in 2003. He is part of the Mobile Platform Architecture and Development team, designing Customer Reference boards for next-generation Intel processors and chipsets. He also works on platform memory power performance initiatives (such as DT in SPD and TS on DIMM) and memory profiling. He has also worked as a research intern at the Indian Space Research Organization. Jayesh received his B.E. degree in Instrumentation & Control Engineering from Gujarat University and his M.S. degree in Electrical Engineering and Computer Engineering from Drexel University, Philadelphia. His e-mail is jayesh.iyer at intel.com.

Corinne Hall joined Intel in 2001 and has worked in both the Desktop Products Organization as well as the Mobile Platform Group. Her current responsibilities include Thermal Sensor on DIMM enabling as well as other memory thermal and power-related activities for DDR2 and DDR3. Corinne received a B.S.E.E. degree in Computer Engineering from Oregon State University. Her e-mail is corinne.l.hall at intel.com.

Jerry Shi has been with Intel for seven years, working as a platform/memory architect. He designed the first Intel® Networking Processor-based cPCI system for demonstration at SuperComm. At Intel he also designed the first set-top-box which could do Picture-in-Picture MPEG playback. He also holds patents related to memory architecture. Prior to Intel, he worked for Philips Semiconductors, Hyundai Electronics, and Oak Technology. Jerry received a B.S.E.E. degree in Computer Science and Engineering from Tsinghua University, China and an M.S.E.E. degree in Electrical Engineering from UMASS at Amherst. His e-mail is jerry.shi at intel.com.

Yuchen Huang is a staff platform memory power Engineer in Intel's Technology Manufacturing Group. She joined Intel in1994, and has held various platform design positions there as Intel's Desktop Product Group platform memory power and performance initiatives owner and as a platform interconnect designer. Currently, Yuchen leads a cross-divisional engineering team within Intel that addresses memory power, thermal, and performance platform challenges. Yuchen received a B.S.E.E. degree in Electrical Engineering from Zhejiang University, China and a M.S.E.E. degree in Electrical Engineering from Portland State University. Her e-mail is yuchen.huang at intel.com.

For further information visit:

developer.intel.com/technology/itj/index.htm